

12 **EUROPEAN PATENT APPLICATION**

21 Application number: 90300372.1

51 Int. Cl.⁵: G06F 13/40, G06F 12/08

22 Date of filing: 12.01.90

30 Priority: 03.02.89 US 306826

43 Date of publication of application:
08.08.90 Bulletin 90/32

84 Designated Contracting States:
AT BE CH DE DK ES FR GB GR IT LI LU NL SE


71 Applicant: **DIGITAL EQUIPMENT CORPORATION**
111 Powdermill Road
Maynard Massachusetts 01754-1418(US)

72 Inventor: **Hetherington, Ricky C.**
177 Indian Meadows
Northboro, Massachusetts 01749(US)
 Inventor: **Fossum, Tryggve**
PO Box 96
Northboro, Massachusetts 01532(US)
 Inventor: **Salett, Ronald M.**
27 Flanagan Drive
Framingham Massachusetts 01701(US)
 Inventor: **Webb, David A. Jr.**
163 Highland Street
Berlin Massachusetts 01503(US)
 Inventor: **Manley, Dwight P.**
45 Gretchen Lane
Holliston Massachusetts 01746(US)

74 Representative: **Maggs, Michael Norman et al**
Kilburn & Strode 30 John Street
GB-London WC1N 2DD(GB)

54 **Method and apparatus for increasing the data storage rate of a computer system.**

57 In a computer system, the flow of data from an execution unit (20) to a cache (28) is enhanced by pairing individual, sequential longword write operations into a simultaneous quadword write operation. A primary and secondary write buffer (50 52) sequentially receive the individual longwords during first and second clock cycles and simultaneously present the individual longwords over a quadword wide bus to the cache (28). During the first clock cycle, when the cache (28) is not performing the quadword write operation, it is free to perform the requisite lookup routine on the address of the first longword of data to determine if the quadword of address space is available in the cache. Thus, the flow of data to the cache 28 is maximised.



METHOD AND APPARATUS FOR INCREASING THE DATA STORAGE RATE OF A COMPUTER SYSTEM

This invention relates generally to an apparatus in a high-speed, digital computer system for controlling the rate at which data is stored and, more particularly, to an apparatus for increasing the data storage bandwidth by combining consecutively located storage requests into a single storage operation.

In one aspect there is provided a method for controlling the flow of data from a longword size bus to a cache of a computer system during a two clock cycle period of time, said cache having a quadword size data access path over which a quadword of data is written to said cache in a single clock cycle, said method comprising the steps of: (a) during a first cycle of said two clock cycle of said two clock cycle period of time, delivering a first longword of data over said bus to said cache and storing said first longword of data in a buffer, and addressing said cache with a cache block address to which said data is to be written to obtain a hit signal when the addressed cache block is free to receive data, and (b) during a second clock cycle of said two clock cycle period of time, in response to said hit signal indicating that the addressed cache block is free to receive data and a quadword context signal indicating that said first longword of data is to be paired with a second longword of data to form a quadword of data and a quadword aligned address signal indicating that the quadword of data is to be stored at an address aligned with said cache block address, delivering said second longword of data over said bus to said cache and storing in said addressed cache block said second longword of data together with said first longword of data from said buffer by passing said first and second longwords of data over respective different portions of said quadword data access path.

This invention will be described with reference to the storage of first and second longwords of data, which are to be stored in a memory as a quadword. It will be appreciated, however, that the invention is equally applicable to data block having other lengths: for example, it could be applicable to first and second blocks of data, each having 64 bits, which are to be stored together in memory as a 1 to 8-bit block.

In the field of high-speed, digital computers it is conventional for a computer system to employ an architecture that is generally of a predefined width, such as 32-bits. Accordingly, most data paths within the computer system are 32-bits wide, including busses, arithmetic logic units, register files, and cache access paths. However, not all data structures within the computer system are of

the same size. In fact, some are narrower, but many are wider, including, for example: double precision floating point numbers; character strings; binary coded decimal strings; 64-bit integers (quadwords); 128-bit integers (octawords); instructions; and stackframes.

These wider data structures are typically employed in high-frequency operations within the computer system. Therefore, in order to increase overall system performance, and prevent bottlenecking, the data paths handling these wider, high-frequency structures have been correspondingly widened. Clearly, by making the data path wider, the amount of data that can be delivered over the path is increased.

There are competing design interests that work against making all data paths wider. First, wider data paths increase the overall cost of the computer system and in some cases offer only negligible increased performance. Alternatively, the wider data path may be needed for only a relative few of its intended operations. Thus, in this case, while the performance increase for individual functions may be dramatic, the overall impact on system performance does not warrant the increased cost.

Finally, while the data structures being communicated may be significantly wider than their data path, the bandwidth of the path may be performance limited, such that simply increasing the path width will have no better effect than optimizing the current data path. For example, in the VAX architecture, the data path from the execution unit to the cache is only 32-bits wide even though the execution unit is capable of performing 64-bit (quadword) storage operations. The quadword is broken down into two 32-bit data structures (longwords) and sequentially transferred over the 32-bit data path. While it may at first seem that the data storage rate could be doubled by increasing the data path to 64-bits, it is not quite that simple. Caching techniques generally require two clock cycles to perform each storage operation. Therefore, even if the data path could deliver 64-bits per cycle, the data storage rate of the cache would only be 64-bits every two cycles.

The present invention is directed to overcoming one or more of the problems as set forth above.

The primary object of the present invention is to increase the rate at which data can be stored in the cache without increasing the width of the data bus connected to the cache.

Another object of the present invention is to provide an apparatus and method for identifying

and pairing consecutive longword storage operations, which are quadword aligned, and storing both longwords in a single storage operation.

In one aspect of the invention, an apparatus is provided for controlling the flow of data to a cache of a computer system. The apparatus includes means for delivering a first longword of data, an address at which the data is to be stored, and a signal indicating that a second longword of data to be stored in the adjacent address will be delivered in the following clock cycle. A primary writebuffer has an output connected to a low-order section of the cache and an input adapted to receive the first longword of data. A secondary writebuffer has an output connected to a high-order section of the cache and an input adapted to receive the second longword of data in response to the first longword being quadword aligned and the second longword of data actually being delivered during the following clock cycle. The apparatus further includes means for substantially simultaneously enabling the high and low-order sections of the cache at the indicated address, whereby the contents of the primary and secondary buffers are stored as a quadword at the address in the cache.

In another aspect of the present invention, a method is provided for controlling the flow of data to a cache of a computer system during a two clock cycle period of time. The method includes the steps of delivering a first longword of data, an address at which the data is to be stored, and a context signal during the first clock cycle. The context signal indicates that a second longword of data to be stored in the adjacent address will be delivered in the second clock cycle. The first longword of data is stored in a primary writebuffer during the first clock cycle. The second longword of data is stored in a secondary writebuffer during the second clock cycle. The second longword is stored in response to the first longword being quadword aligned and the second longword of data actually being delivered during the second clock cycle. The method further including the step of substantially simultaneously enabling the high and low-order sections of the cache at the indicated address during the second clock cycle, whereby the contents of the primary and secondary buffers are stored as a quadword at the address in the cache.

Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the drawings in which:

FIG. 1 is a top level block diagram of a portion of a central processing unit and associated memory;

FIG. 2 is a block diagram of the translation buffer and cache sections of the memory access

unit;

FIG. 3 is a functional diagram of the internal operations of the translation buffer and cache;

FIG. 4 is a timing diagram of significant control events occurring in the translation buffer and memory access unit;

FIG. 5 is a timing diagram of significant control events occurring in the translation buffer and memory access unit during nonoptimized write operations; and

FIG. 6 is a logic diagram of the cache RAM enable signals.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that it is not intended to limit the invention to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

FIG. 1 is a top level block diagram of a portion of a pipelined computer system 10. The system 10 includes at least one central processing unit (CPU) 12 having access to main memory 14. It should be understood that additional CPUs could be used in such a system by sharing the main memory 14. It is practical, for example, for up to four CPUs to operate simultaneously and communicate efficiently through the shared main memory 14.

Inside the CPU 12, the execution of an individual instruction is broken down into multiple smaller tasks. These tasks are performed by dedicated, separate, independent functional units that are optimized for that purpose.

Although each instruction ultimately performs a different operation, many of the smaller tasks into which each instruction is separated are common to all instructions. Generally, the following steps are performed during the execution of an instruction: instruction fetch, instruction decode, operand fetch, execution, and result store. Thus, by the use of dedicated hardware stages, the steps can be overlapped, thereby increasing the total instruction throughput.

The data path through the pipeline includes a respective set of registers for transferring the results of each pipeline stage to the next pipeline stage. These transfer registers are clocked in response to a common system clock. For example, during a first clock cycle, the first instruction is fetched by hardware dedicated to instruction fetch. During the second clock cycle, the fetched instruction is transferred and decoded by instruction decode hardware, but, at the same time, the next instruction is fetched by the instruction fetch hard-

ware. During the third clock cycles each instruction is shifted to the next stage of the pipeline and a new instruction is fetched. Thus, after the pipeline is filled, an instruction will be completely executed at the end of each clock cycle.

This process can be analogized to an assembly line in a manufacturing environment. Each worker is dedicated to performing a single task on every product that passes through his or her work stage. As each task is performed, the product comes closer to completion. At the final stage, each time the worker performs his assigned task a completed product rolls off the assembly line.

As shown in FIG. 1, the CPU 12 is partitioned into at least three functional units: a memory access unit 16, an instruction unit 18, and an execution unit 20. These units are sometimes referred to as the MBOX, IBOX and EBOX, respectively.

The instruction unit 18 prefetches instructions, decodes opcodes to obtain operand and result specifiers, fetches operands, and updates a program counter 24. The instruction unit 18 includes an operand processing unit 22, the program counter 24, and an instruction decoder 26. The program counter 24 is maintained in the instruction unit 18 so that the proper instructions can be retrieved from a high-speed cache memory 28 maintained in the memory access unit 16. The cache 28 stores a copy of a small portion of the information stored in main memory 14 and is employed to increase processing speed by reducing memory access time. Operation of the cache 28 is described in greater detail below in conjunction with the description of the memory access unit 16.

The program counter 24 preferably uses virtual memory locations rather than the physical memory locations of the main memory 14 and cache 28. Thus, the virtual address of the program counter 24 must be translated into the physical address of the main memory 14 before instructions can be retrieved. Accordingly, the contents of the program counter 24 are transferred to the memory access unit 16 where a translation buffer 30 performs the address conversion. The instruction is retrieved from its physical memory location in the cache 28 using the converted address. The cache 28 delivers the instruction over the data return lines 32 to the instruction decoder 26. The organization and operation of the cache 28 and translation buffer 30 are further described in Chapter 11 of Levy and Eckhouse, Jr., Computer Programming and Architecture, The VAX-11, Digital Equipment Corporation, pp. 351-368 (1980).

The operand processing unit (OPU) 22 also produces virtual addresses. In particular, the OPU 22 produces virtual addresses for memory source (read) and destination (write) instructions. For at least the memory read instructions, the OPU 22

must deliver these virtual addresses to the memory access unit 16 where they are translated to physical addresses. The physical memory locations of the cache 28 are then accessed to fetch the operands for the memory source instructions.

In order to practice the preferred method of the present invention, the OPU 22 also delivers to the memory access unit 16 the virtual addresses of the destinations for the memory destination instruction operands. The virtual address, for example, is a 32-bit number. In addition to transmitting the 32-bit virtual address, the OPU 22 also delivers a 3-bit control field to indicate whether the instruction specifies a read or write operation. In the event that the control field indicates that the virtual address corresponds to a read instruction, the cache 28 retrieves the data from the identified physical memory location and delivers it over data return lines 34 to the execution unit 20.

Conversely, for a write operation the write address is stored until the data to be written is available. Clearly, for instructions such as MOVE or ADD, the data to be written is not available until execution of the instruction has been completed. However, the virtual address of the destination can be translated to a corresponding physical address during the time required for execution of the instruction. Also, it is desirable for the OPU 22 to preprocess multiple instruction specifiers during this time in order to increase the overall rate at which instructions are performed. For these purposes, the memory access unit 16 is provided with a "write queue" 36 intermediate the translation buffer 30 and the cache 28 for storing the physical destination addresses of a variable number of write operations. The write queue 36 maintains the address until the execution unit 20 completes the instruction and sends the resulting data to the memory access unit 16. This data is paired with the previously stored write address and written into the cache 28 at that memory location.

The OPU 22 also operates on instructions which are not memory operands. For example, the OPU 22 also processes immediate operands, short literals and register operands. In each of these types of instructions, the OPU 22 delivers its results directly to the execution unit 20.

The first step in processing the instructions is to decode the "opcode" portions of the instruction. The first portion of each instruction consists of its opcode which specifies the operation to be performed in the instruction. The decoding is done using a standard table-look-up technique in the instruction decoder 26. The instruction decoder 26 finds a microcode starting address for executing the instruction in a look-up table and passes that starting address to the execution unit 20. Later, the execution unit 20 performs the specified operation

by executing prestored microcode, beginning at the indicated starting address. Also, the decoder 26 determines where source-operand and destination-operand specifiers occur in the instruction and passes these specifiers to the operand processing unit 22 for preprocessing prior to execution of the instruction.

Referring now to FIG. 2, the memory access unit 16 includes the cache 28, the translation buffer 30, the write queue 36, and a group of registers 38. As noted above, the cache 28 is a high-speed memory that stores a copy of a small portion of the information stored in the main memory 14. The cache 28 is accessible at a much higher rate than the main memory 14. Its purpose, therefore, is to reduce the average time necessary for a memory access (i.e., a read or write) to be performed. Since the cache 28 stores only a small portion of the information stored in the main memory 14, there will occasionally be instructions that attempt to access memory not contained in the cache 28. The cache 28 recognizes when these "misses" occur, and in these instances the cache 28 retrieves the identified data from the main memory 14. Of course, during these "misses" performance of the CPU 12 will suffer. However, with the cache 28 the overall memory access speed is increased.

The translation buffer 30 is a high-speed, associative memory that stores the most recently used virtual-to-physical address translations. In a virtual memory system, a reference to a single virtual address can cause several memory references before the desired information is made available. However, where the translation buffer 30 is used, translation is reduced to simply finding a "hit" in the translation buffer 30. These virtual addresses generated by the OPU 22 and execution unit 20 are stored in latches 35, where they are maintained until they are accessed via a multiplexer 37 and serviced by the translation buffer 30.

Once the virtual-to-physical address translation is complete, the physical address is transferred to the write queue 36 or one of the registers 38. As its name suggests, the write queue 36 receives the physical address only if the corresponding operation is a write to memory. The purpose of the write queue 36 is to provide a temporary storage location for the physical write address of the write operation. Because of the pipeline nature of the CPU 12, the write address is available before the data to be stored therein is available. In fact, the data will only become available after the execution of the instruction in the execution unit 20. Moreover, because it is desired to preprocess multiple operand specifiers for instructions in the pipeline, it is likely that there will be a plurality of physical write addresses waiting for their corresponding data. Accordingly, the write queue 36 is a multiple

position first-in, first-out buffer constructed to accommodate a plurality of physical write addresses.

Conversely, if the operation corresponding to the physical address is a read operation, then the translation buffer 30 provides the physical address for an operand of the read operation. The read address is transferred to one of the registers 38 where it is selected by a multiplexer 40 and delivered to the cache 28. The cache 28 accesses the identified memory location and delivers the data stored at that location to the execution unit 20 via the data return lines 34.

The cache 28 is divided into two sections, a data storage area and a tag storage area. Since the cache 28 contains only a portion of the main memory 14, the tag storage area is necessary in order to keep track of what data is currently located in the data storage area. Thus, during a cache read operation, the tag and data storage areas are accessed in the same clock cycle, using the physical address in one of the registers 38. If the desired data is available in the cache, then read data is immediately available in the next clock cycle. As long as the requested data is available in the cache 28, then the cache 28 is capable of performing one read operation every clock cycle.

Conversely, the cache 28 is only capable of performing one write operation every other clock cycle. During a write operation, the tag storage area must be interrogated before the new data is written. Otherwise, data already present in the cache 28 could be overwritten and destroyed. Thus, during a write operation, the tag storage area is accessed in the first clock cycle and the data storage area is accessed in the second clock cycle.

Accordingly, it can be seen that even if the data path between the execution unit 20 and the cache 28 is 64-bits wide, data is stored in the cache 28 at the rate of 32-bits per clock cycle (64-bits every two cycles). The maximum bandwidth of the data path is 32-bits per cycle. However, since the cache 28 is capable of performing 64-bit storage operations, then the 32-bit data path can perform at the same rate as a 64-bit data path if consecutive 32-bit write operations can be paired together and stored in one 64-bit operation.

Where multiple 32-bit words are to be written, they are usually adjacent in memory. Furthermore, most data is naturally aligned in memory. A naturally aligned quadword has an address in which the three least significant bits are zero. Also note that a quadword is composed of two longwords. It is, therefore, likely that two consecutive longword write operations from the execution unit 20 will fit within the same aligned quadword in the cache 28. This is typically true for double precision floating point data, string data, procedure call stack frames, etc.

A primary and secondary writebuffer 50, 52 are connected in parallel to the 32-bit data bus from the execution unit 20. The outputs of the primary and secondary write buffers 50, 52 are respectively connected to the lower and upper 32-bits of a 64 bit data path into the cache 28. This 64-bit data path also interconnects the cache 28 with the main memory 14 and is used to refill the cache 28 from the main memory 14. It should be noted that during a cache refill the tag storage area does not need to be interrogated. Unlike an execution unit initiated write operation, a write operation during a cache refill can be performed during every clock cycle. Thus, in this case the 64-bit data path optimizes the data storage rate for cache refills.

The primary writebuffer 50 ordinarily receives and stores the data to be written to the cache 28. The secondary writebuffer 52 only receives and stores data in the event that the execution unit 20 delivers the second of two consecutive longword write operations. Thereafter, the cache 28 accepts a longword of data from each of the writebuffers 50, 52. The secondary writebuffer 52 is only used during this optimization of paired longword write operations.

Referring now to FIG. 3, a functional diagram of the internal operation of the translation buffer 30, cache 28, and writebuffer 50, 52 control signals is shown. The translation buffer 30 receives four different type signals from the execution unit 20: a 32-bit virtual address; a 1-bit address valid signal; a 5-bit command signal; and a 3-bit context signal. The 32-bit virtual address is, as discussed above, stored in the latch 35 from where it is ultimately accessed by the multiplexer 37 and converted from a virtual to a physical address. The low-order bits act as a pointer into the RAM 56. The high-order address bits of the data actually stored in that RAM location are presented to a comparator 58 along with the high-order bits of the virtual address. If they match, then the address stored in the RAM location is the corresponding physical address and it is clocked into a buffer 60 by the output signal of the comparator 58.

At the same time, in order to determine if this address corresponds to the first longword of an optimized quadword write operation, the translation buffer 30 must determine if three conditions are satisfied. First, the address must be quadword aligned in order to perform a quadword write operation. To determine if the address is quadword aligned, it is only necessary to inspect the low-order 3-bits of the virtual address. Accordingly, a 3-bit comparator 62 has a first input connected to the low-order 3-bits of the virtual address and a second input connected to a preselected constant value of 000.

The second condition requires that the execu-

tion unit 20 actually be requesting a quadword write operation. The 3-bit context signal provided by the execution unit 20 contains a preselected code that identifies the size of operation to be performed while the 5-bit command field indicates the type (i.e. write) operation. The execution unit 20 can request quadword, longword, or byte write operations. The optimization will only occur if the execution unit 20 has requested a quadword write operation. Thus, a 3-bit comparator 64 has a first input connected to the context signal and a second input connected to a preselected constant value that matches the code for a quadword write request.

The outputs of the comparators 62, 64 are connected to the inputs of a 3-input AND gate 66. The third input to the AND gate 66 is connected directly to the address valid signal from the execution unit 20. The address valid signal indicates that the execution unit 20 has properly delivered the subsequent longword address and corresponding data in time for the quadword optimization to occur. Thus, the AND gate 66 delivers a 1-bit quadword valid signal to the cache 28, thereby enabling the cache 28 to receive a longword of data from each of the writebuffers 50, 52.

Within the cache 28, a buffer control 68 receives the quadword valid signal along with a data valid signal from the execution unit 20. The data valid signal is delivered by the execution unit 20 to indicate that data has been placed on the 32-bit data bus. Ordinarily, during nonoptimized data transfers and during the transfer of the low-order longword of an optimized data transfer, the buffer control 68 produces a hold signal to the primary writebuffer 50, allowing the writebuffer 50 to store the data currently presented on the data bus. The buffer control 68 produces this primary hold signal in response to receiving the data valid signal in the absence of the quadword valid signal.

On the other hand, when both the quadword and data valid signals are present, the buffer control 68 outputs a hold signal to the secondary writebuffer 52, causing it to store the data currently present on the bus. In this manner, during an optimized quadword write operation, the writebuffers 50, 52 are consecutively loaded with the lower and upper longwords of the quadword data.

A 64-bit RAM array 70 contained within the cache 28 is divided into two sets of 32-bit storage locations where each 32-bit set has an independently operable enable input. The 32-bit physical address from the translation buffer 30 acts as a pointer into the RAM array 70 and both of the enable inputs are connected to the quadword valid signal. Thus, during an optimized quadword write when the quadword valid signal is asserted, both longwords of the RAM array 70 are enabled to

store the two longwords currently held in the writebuffers 50, 52.

The timing and operation of the optimized quadword write operation may be better appreciated by referring to the timing diagrams illustrated in FIG. 4. Five clock cycle periods A-E are illustrated for the context, data, and quadword valid signals, as well as the cache lookup and write operations. Beginning in clock cycle A, the execution unit 20 issues a context signal indicating that the data and address currently being delivered corresponds to the first longword of a quadword write operation. Shortly thereafter, the first 32-bits of data along with the data valid signal are delivered from the execution unit 20 to the primary writebuffer 50. In the absence of the quadword valid signal, the buffer control 68 enables the primary writebuffer 50 to save the first longword of data. At the same time, the translation buffer 30 performs the virtual to physical address conversion and delivers the resulting physical address to the cache 28.

In clock cycle B, the execution unit 20 issues another context signal indicating that the data and address currently being delivered correspond to the second longword of a quadword write operation. Shortly thereafter, the second 32-bits of data along with the data valid signal are delivered from the execution unit 20 to the secondary writebuffer 52. Since the execution unit 20 has successfully delivered the second longword of data, the address valid signal is asserted, thereby causing the quadword valid signal to be similarly asserted. The presence of the quadword valid signal results in the buffer control 68 enabling the primary writebuffer 52 to save the second longword of data.

It should be remembered that the cache 28 requires two clock cycles to perform a write operation. The first clock cycle involves looking up the tag in the cache ram array 70 to prevent overwriting good data and the second clock cycle is dedicated to actually writing the data into the ram array 70. Thus, during clock cycle B the cache 28 performs the lookup function.

In clock cycle C, the asserted quadword valid signal ensures that both 32-bit sections of the ram array 70 are enabled so that both the primary and secondary writebuffers 50, 52 are loaded into the ram array 70. At the same time, the execution unit 20 is sending the quadword context signal for the next quadword of data to be written into the cache 28. Clock cycles C and D are substantially identical to clock cycles A and B. Thus, during clock cycle C and every second clock cycle thereafter, a 64-bit cache write operation is performed.

It should be appreciated that the logical conditions described in the timing diagrams of FIG. 4 represent the maximum data transfer rate from the

execution unit 20 to the cache 28. Thus, using only a 32-bit data bus and constrained by the cache 28 being able to perform only one write operation every two clock cycles, the instant invention achieves an effective transfer rate of 32-bits per clock cycle.

In contradistinction thereto, the timing diagrams of FIG. 5 represent the transfer of data between the execution unit 20 and the cache 28 where an optimized quadword transfer is attempted, but fails. Even though the optimized quadword transfer fails, an ordinary longword transfer is still accomplished, allowing the CPU to continue operating, albeit at a temporarily slower rate.

In clock cycle A, the execution unit 20 issues the quadword address valid, indicating that the data and address currently being delivered correspond to the first longword of a quadword write operation. Shortly thereafter, the first 32-bits of data along with the data valid signal are delivered from the execution unit 20 to the primary writebuffer 50. In the absence of the quadword valid signal, the buffer control 68 enables the primary writebuffer 50 to save the first longword of data. At the same time, the translation buffer 30 performs the virtual to physical address conversion and delivers the resulting physical address to the cache 28.

In clock cycle B, the execution unit 20 fails to issue another context signal, thereby indicating that the desired longword data and address are not currently being delivered. Thus, the quadword valid signal is not asserted, the secondary writebuffer 52 is not enabled to save any data present on the bus, and both 32-bit sections of the ram array 70 are not enabled. The cache 28 performs the lookup in clock cycle B and the write operation in clock cycle C, but only the lower 32-bit section of the ram array 70 is enabled to receive only the contents of the primary writebuffer 50.

Therefore, the effective data transfer rate is only one half the optimized quadword transfer rate. Here, a 32-bit longword is transferred every second clock cycle thereafter.

Further, it should be noted that the timing diagram for a failed optimized quadword transfer is substantially identical to an ordinary longword transfer. The only difference is in clock cycle A where the execution unit 20 delivers a context signal corresponding to a longword transfer rather than a quadword transfer. Therefore, even a failed optimized quadword transfer has the same effective transfer rate as an unoptimized longword transfer.

FIG. 6 is a logic diagram of the cache RAM enable signals. The AND gate 66 receives inputs from a pair of latches 80, 82 and the address valid signal from the execution unit 20. The latch 80 has an input connected to the output of the comparator

62 where that signal is indicative of the quadword address being quadword aligned. The latch 82 has an input connected to the output of the comparator 64 where that signal is representative of a context signal from the execution unit 20 indicating a quadword write is being requested. The output of the AND gate 66 is the quadword valid signal, which is passed through a pair of OR gates 84, 86 to the enable inputs of the high and low 32-bit sections of the cache ram array 70.

Each of the OR gates 82, 84 also has a second input for enabling the high and low 32-bit sections of the ram array 70. The high 32-bit section is also enabled when the context signal corresponds to a longword write request and the address of the longword write request corresponds to the high 32-bit section of the ram array 70. An AND gate 88 receives inputs of context=longword and address=001. Thus, the output of the AND gate 88 is asserted only when the execution unit has requested a longword write operation and the address to be written corresponds to the upper 32-bit section.

Similarly, the low 32-bit section is also enabled when the context signal corresponds to a longword write request and the address of the longword write request corresponds to the low 32-bit section of the ram array 70. An AND gate 90 receives inputs of context=longword and address=000. Further, the original context signal requesting a quadword transfer is demoted to a longword request and passed to the inputs of the AND gates 88, 90.

Therefore, when the optimized quadword transfer is possible, the quadword valid signal is passed through the OR gates 84, 86 to the high and low enable inputs of the RAM array 70. On the other hand, even where the optimized quadword transfer fails, the demoted quadword context signal is passed through the appropriate AND gate 88, 90 to either the low or high enable inputs of the RAM array 70.

Claims

1. An apparatus for controlling the flow of data to a cache of a computer system, comprising:
means for delivering a first longword of data, an address at which the data is to be stored, and a signal indicating that a second longword of data to be stored in the adjacent address will be delivered in the following clock cycle;
a primary writebuffer having an output connected to a low-order section of the cache and an input adapted to receive the first longword of data;
a secondary writebuffer having an output connected to a high-order section of the cache and an input adapted to receive the second longword of

data in response to said first longword being quadword aligned and said second longword of data actually being delivered during said following clock cycle; and

means for substantially simultaneously enabling the high and low order sections in said cache at said address, whereby the contents of the primary and secondary buffers are stored as a quadword at said address in the cache.

2. An apparatus, as set forth in claim 1, wherein said primary and secondary writebuffers are each configured to be one longword wide and there inputs are connected in parallel to a data bus from an execution unit of the computer system.

3. An apparatus, as set forth in claim 1, wherein said cache is adapted to perform a lookup to determine if said quadword aligned address is available in the cache, said lookup being performed during the same clock cycle in which said address is received whereby said cache is free to store said quadword of data immediately after said second longword of data is received by said secondary writebuffer.

4. An apparatus, as set forth in claim 1, including means for preventing the simultaneous enabling of the high and low-order sections of the address in said cache in response to the absence of said second longword of data in the following clock cycle.

5. An apparatus, as set forth in claim 4, including means for enabling only the low-order section of the address in the cache in response to the absence of the second longword of data in the following clock cycle.

6. A method for controlling the flow of data to a cache of a computer system during a two clock cycle period of time, comprising the steps of:
delivering a first longword of data, an address at which the data is to be stored, and a context signal during the first clock cycle, said context signal indicating that a second longword of data to be stored in the adjacent address will be delivered in the second clock cycle;

storing the first longword of data in a primary writebuffer during the first clock cycle;
storing the second longword of data in a secondary writebuffer during the second clock cycle, said second longword being stored in response to said first longword being quadword aligned and said second longword of data actually being delivered during said second clock cycle; and

substantially simultaneously enabling the high and low-order sections in said cache at said address during the second clock cycle, whereby the contents of the primary and secondary buffers are stored as a quadword at said address in the cache.

7. A method, as set forth in claim 6, including the step of preventing the simultaneous enabling of

the high and low-order sections of the address in said cache in response to the absence of said second longword of data in the second clock cycle.

8. A method, as set forth in claim 7, including the step of enabling only the low-order section of the address in the cache during the second clock cycle in response to the absence of the second longword of data in the second clock cycle.

9. A method, as set forth in claim 8, including the step of looking up the quadword aligned address in the cache during the first clock cycle to determine that said address is available for storing the quadword of data.

5

10

15

20

25

30

35


40


45

50

55

FIG 1





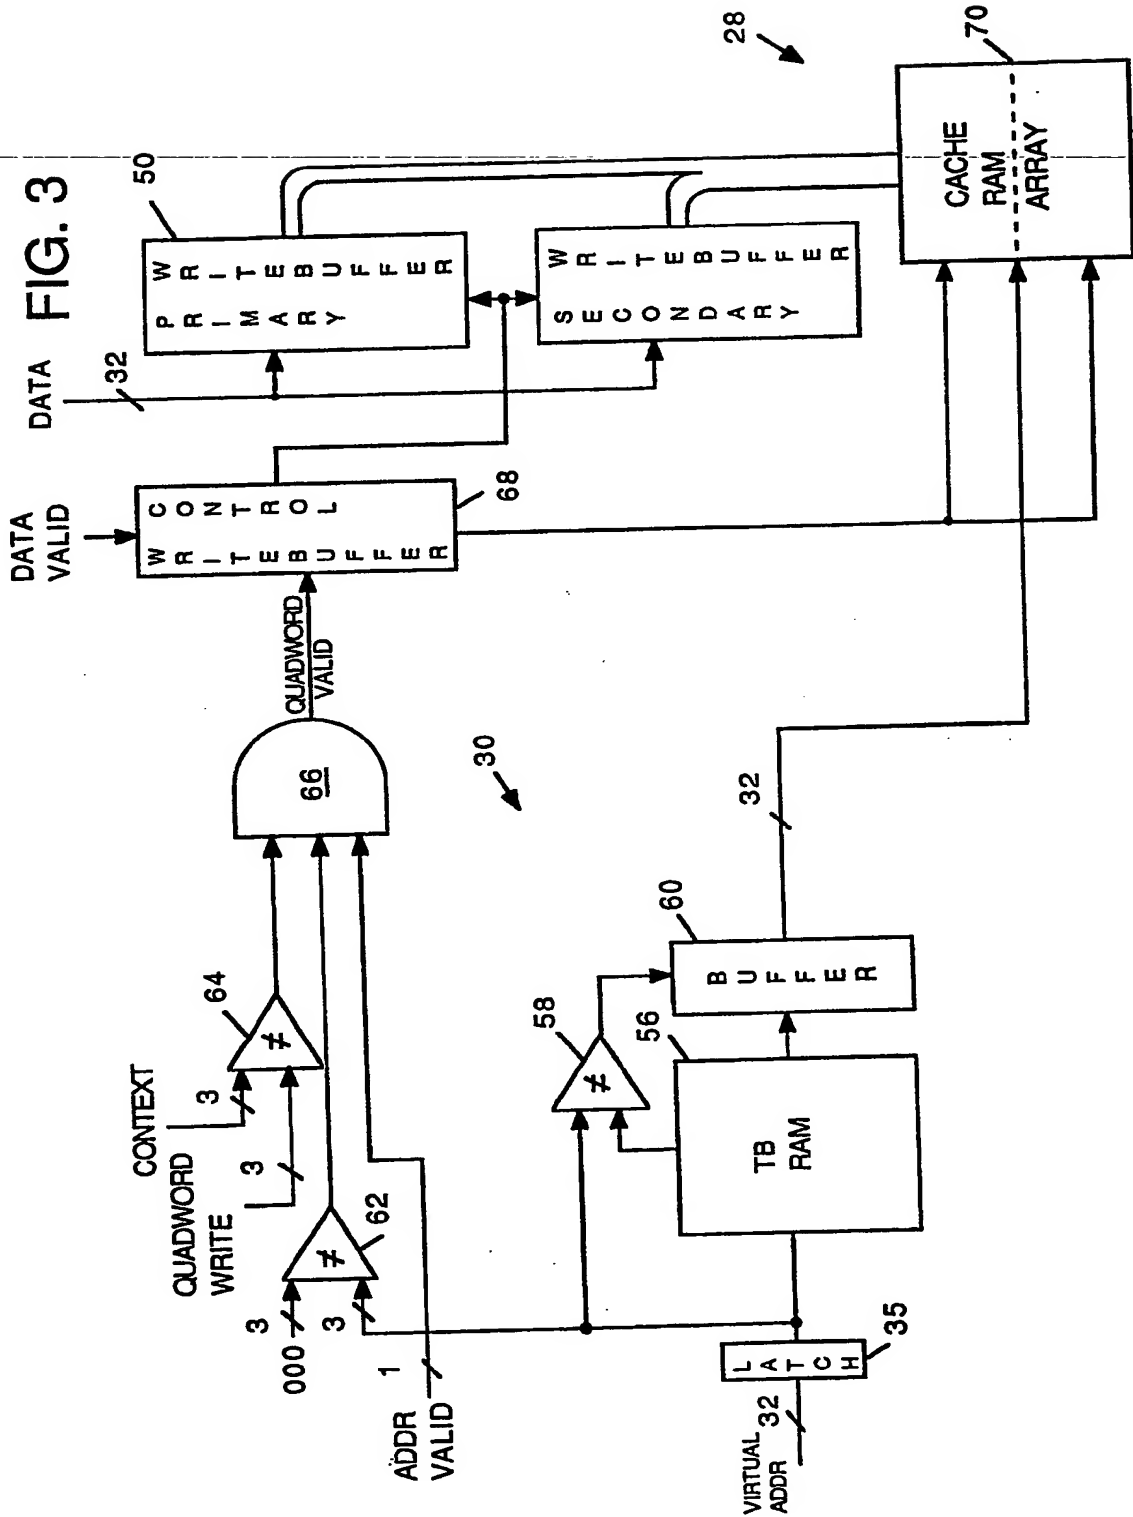


FIG. 4




FIG. 5





FIG. 6



(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 381 323 A3

(12)


EUROPEAN PATENT APPLICATION(21) Application number: **90300372.1**(51) Int. Cl.⁵: **G06F 12/08, G06F 13/40**(22) Date of filing: **12.01.90**(30) Priority: **03.02.89 US 306826**(43) Date of publication of application:
08.08.90 Bulletin 90/32(84) Designated Contracting States:
AT BE CH DE DK ES FR GB GR IT LI LU NL SE(88) Date of deferred publication of the search report:
06.05.92 Bulletin 92/19(71) Applicant: **DIGITAL EQUIPMENT CORPORATION**
111 Powdermill Road
Maynard Massachusetts 01754-1418(US)(72) Inventor: **Hetherington, Ricky C.**
177 Indian Meadows
Northboro, Massachusettes 01749(US)
Inventor: **Fossum, Tryggve**
PO Box 96
Northboro, Massachusettes 01532(US)
Inventor: **Salett, Ronald M.**
27 Flanagan Drive
Framingham Massachusettes 01701(US)
Inventor: **Webb, David A. Jr.**
163 Highland Street
Berlin Massachusettes 01503(US)
Inventor: **Manley, Dwight P.**
45 Gretchen Lane
Holliston Massachusettes 01746(US)(74) Representative: **Maggs, Michael Norman et al**
Kilburn & Strode 30 John Street
GB-London WC1N 2DD(GB)(54) **Method and apparatus for increasing the data storage rate of a computer system.**

(57) In a computer system, the flow of data from an execution unit (20) to a cache (28) is enhanced by pairing individual, sequential longword write operations into a simultaneous quadword write operation. A primary and secondary write buffer (50 52) sequentially receive the individual longwords during first and second clock cycles and simultaneously present the individual longwords over a quadword

wide bus to the cache (28). During the first clock cycle, when the cache (28) is not performing the quadword write operation, it is free to perform the requisite lookup routine on the address of the first longword of data to determine if the quadword of address space is available in the cache. Thus, the flow of data to the cache 28 is maximised.

EP 0 381 323 A3

FIG. 2





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 90 30 0372

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
A	US-A-4 381 541 (BAUMANN) * column 3, line 21 - column 4, line 13 * * column 6, line 19 - line 32 * * column 9, line 61 - column 10, line 19 * * figure 4 * ---	1,2,6	G06F12/08 G06F13/40
A	US-A-4 317 168 (MESSINA) * column 2, line 5 - line 10 * * column 3, line 51 - line 66 * * column 4, line 60 - line 68 * * column 5, line 31 - line 40 * ---	1,6	
A	WO-A-8 504 737 (AT&T) * page 26, line 2 - line 25 * * figure 6 * -----	1,6	
			TECHNICAL FIELDS SEARCHED (Int. Cl.5)
			G06F
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 26 FEBRUARY 1992	Examiner MASCHE C.
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons ----- & : number of the same patent family, corresponding document			

THIS PAGE BLANK (USPTO)